



US009043331B2

(12) **United States Patent**  
**Viridy**

(10) **Patent No.:** **US 9,043,331 B2**

(45) **Date of Patent:** **May 26, 2015**

(54) **SYSTEM AND METHOD FOR INDEXING DOCUMENTS ON THE WORLD-WIDE WEB**

(75) Inventor: **Ajaipal Singh Viridy**, Loudoun County, VA (US)

(73) Assignee: **FACEBOOK, INC.**, Menlo Park, CA (US)

(\* ) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **13/620,911**

(22) Filed: **Sep. 15, 2012**

(65) **Prior Publication Data**

US 2013/0173629 A1 Jul. 4, 2013

**Related U.S. Application Data**

(63) Continuation of application No. 12/685,799, filed on Jan. 12, 2010, now Pat. No. 8,768,912, which is a continuation of application No. 12/054,130, filed on Mar. 24, 2008, now Pat. No. 7,664,767, which is a

(Continued)

(51) **Int. Cl.**

**G06F 17/30** (2006.01)

**G06Q 30/02** (2012.01)

(52) **U.S. Cl.**

CPC .... **G06F 17/30864** (2013.01); **G06F 17/30616** (2013.01); **G06F 17/30705** (2013.01); **G06F 17/3071** (2013.01); **G06Q 30/02** (2013.01); **Y10S 707/99943** (2013.01); **Y10S 707/99933** (2013.01); **Y10S 707/99945** (2013.01); **Y10S 707/99942** (2013.01); **Y10S 707/99944** (2013.01); **G06F 17/30241** (2013.01); **G06F 17/30613** (2013.01); **G06F 17/30321** (2013.01)

(58) **Field of Classification Search**

None

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,255,796 A 3/1981 Gabbe et al.

4,586,091 A 4/1986 Panaoussis

4,809,081 A 2/1989 Linehan

(Continued)

FOREIGN PATENT DOCUMENTS

GB 2114407 8/1983

WO WO 93/18484 9/1993

(Continued)

OTHER PUBLICATIONS

Williamson, S.; Transition and Modernization of the Internet Registration Service; Mar. 10, 1993; Network Solutions, Inc.; RFC 1400.\*

(Continued)

*Primary Examiner* — Tarek Chbouki

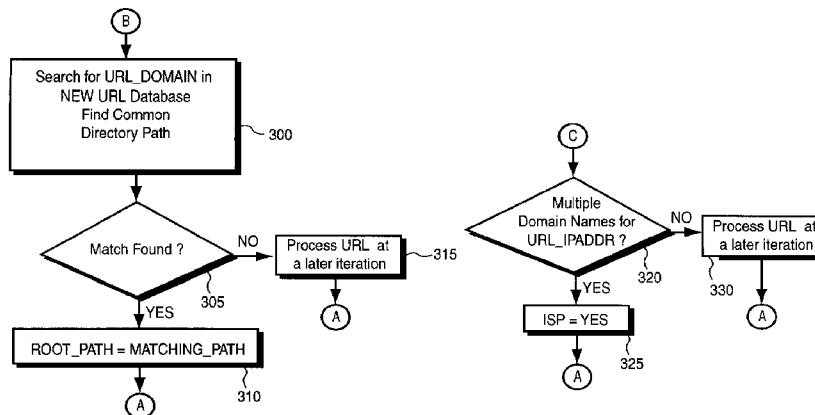
*Assistant Examiner* — Michael K Tamaru

(74) *Attorney, Agent, or Firm* — Keller Jolley Preece

(57) **ABSTRACT**

A method and search engine for classifying a source publishing a document on a portion of a network, includes steps of electronically receiving a document, based on the document, determining a source which published the document, and assigning a code to the document based on whether data associated with the document published by the source matches with data contained in a database. An intelligent geographic- and business topic-specific resource discovery system facilitates local commerce on the World-Wide Web and also reduces search time by accurately isolating information for end-users. Distinguishing and classifying business pages on the Web by business categories using Standard Industrial Classification (SIC) codes is achieved through an automatic iterative process.

**38 Claims, 6 Drawing Sheets**



**Related U.S. Application Data**

continuation of application No. 10/705,822, filed on Nov. 13, 2003, now Pat. No. 7,349,892, which is a continuation of application No. 09/500,471, filed on Feb. 9, 2000, now Pat. No. 6,691,105, which is a continuation of application No. 08/844,522, filed on Apr. 18, 1997, now Pat. No. 6,148,289.

(60) Provisional application No. 60/017,548, filed on May 10, 1996.

**References Cited**

U.S. PATENT DOCUMENTS

(56)

5,003,614	A	3/1991	Tanaka et al.	
5,200,993	A	4/1993	Wheeler et al.	
5,206,949	A	4/1993	Cochran et al.	
5,233,513	A	8/1993	Doyle	
5,307,266	A	4/1994	Hayashi et al.	
5,365,434	A	11/1994	Figliuzzi	
5,384,835	A	1/1995	Wheeler et al.	
5,404,510	A	4/1995	Smith et al.	
5,412,804	A	5/1995	Krishna	
5,444,779	A	8/1995	Daniele	
5,452,445	A	9/1995	Hallmark et al.	
5,485,608	A	1/1996	Lomet et al.	
5,485,610	A	1/1996	Gioielli et al.	
5,495,608	A	2/1996	Antoshenkov	
5,500,929	A	3/1996	Dickinson	
5,510,778	A	4/1996	Krieter et al.	
5,513,314	A	4/1996	Kandasamy et al.	
5,530,852	A	6/1996	Meske, Jr. et al.	
5,544,049	A	8/1996	Henderson et al.	
5,572,643	A	11/1996	Judson	
5,579,407	A	11/1996	Murez	
5,600,831	A	2/1997	Levy et al.	
5,608,907	A *	3/1997	Fehskens et al. ....	718/102
5,640,553	A	6/1997	Schultz	
5,724,517	A *	3/1998	Cook et al. ....	709/227
5,737,533	A	4/1998	De Hond	
5,751,961	A	5/1998	Smyk	
5,764,906	A	6/1998	Edelstein et al.	
5,778,364	A	7/1998	Nelson	
5,778,367	A	7/1998	Wesinger, Jr. et al.	
5,794,216	A	8/1998	Brown	
5,805,810	A	9/1998	Maxwell	
5,826,025	A	10/1998	Gramlich	
5,828,990	A	10/1998	Nishino et al.	
5,842,174	A	11/1998	Yanor	
5,855,020	A	12/1998	Kirsch	
5,870,550	A	2/1999	Wesinger et al.	
5,878,233	A	3/1999	Schloss	
5,878,398	A	3/1999	Tokuda et al.	
5,926,812	A	7/1999	Hilsenrath et al.	
5,930,474	A	7/1999	Dunworth et al.	
5,948,040	A	9/1999	DeLorme et al.	
5,956,716	A	9/1999	Kenner et al.	
6,002,853	A	12/1999	De Hond	
6,035,330	A *	3/2000	Astiz et al. ....	709/218
6,148,289	A	11/2000	Viridy	
6,243,724	B1	6/2001	Mander	
6,577,755	B1	6/2003	Lorie	
6,691,105	B1	2/2004	Viridy	
6,999,574	B2	2/2006	Baker	
7,139,252	B2	11/2006	Babu et al.	
7,349,892	B1	3/2008	Viridy	
7,664,767	B2	2/2010	Viridy	
8,244,577	B2	8/2012	Gupta	
2003/0130993	A1	7/2003	Mendelevitch et al.	
2003/0231785	A1	12/2003	Rhoads et al.	
2004/0161097	A1	8/2004	Henry	
2005/0033777	A1	2/2005	Moraes et al.	
2006/0059129	A1	3/2006	Azuma et al.	
2006/0245641	A1	11/2006	Viola et al.	
2006/0253324	A1	11/2006	Miller	
2007/0083509	A1	4/2007	Wu et al.	

2007/0208669	A1	9/2007	Rivette et al.
2007/0286169	A1	12/2007	Roman
2008/0071752	A1	3/2008	Azuma et al.
2008/0130945	A1	6/2008	Rhoads et al.
2008/0228574	A1	9/2008	Stewart et al.
2008/0255977	A1	10/2008	Altberg et al.
2009/0070290	A1	3/2009	Nye
2009/0080635	A1	3/2009	Altberg et al.
2010/0114904	A1	5/2010	Viridy
2010/0287282	A1	11/2010	MacLaughlin
2012/0278361	A1	11/2012	Gupta
2013/0110817	A1	5/2013	Viridy
2013/0173624	A1	7/2013	Viridy
2013/0173625	A1	7/2013	Viridy
2013/0173626	A1	7/2013	Viridy
2013/0173628	A1	7/2013	Viridy
2013/0173630	A1	7/2013	Viridy
2013/0191367	A1	7/2013	Viridy

FOREIGN PATENT DOCUMENTS

WO	WO 95/08809	3/1995
WO	WO 95/09395	4/1995

OTHER PUBLICATIONS

Database 16 (IAC PROMT) on Dialog, No. 6303723, "Imperative! Announces New Site for Locating Companies on the Internet", PR Newswire, 1 page, Jul. 22, 1996.

Database 148 (Trade and Industry Database) on Dialog, No. 7559252, "Open Market Inc. offers Internet Users Free Access to Directory of Commerical Sites on the Internet", Business Wire, 2 pages, Nov. 8, 1994.

Schwartz et al., "Applying an Information Gathering Architecture to Netfind: A White Pages Tool for a Changing and Growing Internet" IEE/ACM Transactions on Networking, vol. 2, No. 5, Oct. 1994, pp. 426-439.

Office Action in U.S. Appl. No. 10/705,822, dated Aug. 30, 2006, 14 pages.

Final Office Action in U.S. Appl. No. 10/705,822 dated Feb. 1, 2007, 14 pages.

Office Action in U.S. Appl. No. 10/705,822, dated Jul. 17, 2007, 12 pages.

Notice of Allowance in U.S. Appl. No. 10/705,822 dated Nov. 6, 2007, 29 pages.

Office Action in U.S. Appl. No. 12/054,130, dated May 7, 2009, 32 pages.

Notice of Allowance in U.S. Appl. No. 12/054,130 dated Sep. 28, 2009, 16 pages.

Esposito et al., Empirical Learning Methods for Digitized Document Recognition: An Integrated Approach to Inductive Generalization, IEEE 1990, pp. 37-45.

Esposito et al., An Experimental Page Layout Recognition System for Office Document Automatic Classification: An Integrated Approach for Inductive Generalization, IEEE 1990, pp. 557-562.

Lin et al., An Automatic Indexing and Neutral Networking Approach to Concept Retrieval and Classification of Multilingual Documents, IEEE Transactions on systems, man and cybernetics—part B: Cybernetics, vol. 26, No. 1, Feb. 1996, pp. 75-88.

Mitchell et al., Automated Categorization of Real-Time Newswire Stories, Proceedings of the Twenty-Ninth Hawaii International Conference on System Sciences, Jan. 1996, pp. 92-97.

Perkowitz et al., Category Translation: Learning to understand information on the Internet, AAAI Technical Reports SS-95-08, 1995, pp. 142-148.

Summers, Toward a taxonomy of logical document structures, 1995, Proceedings of the Dartmouth Institute for Advanced Graduate Studies.

Conway, Page Grammars and Page Parsing A Syntactic Approach to Document Layout Recognition, Oct. 1993, Document Analysis and Recognition, 1993., Proceedings of the Second International Conference.

U.S. Appl. No. 12/685,799, Apr. 29, 2013, Office Action.

U.S. Appl. No. 13/612,390, Jun. 26, 2013, Office Action.

(56)

**References Cited**

OTHER PUBLICATIONS

U.S. Appl. No. 13/612,512, Aug. 2, 2013, Office Action.  
U.S. Appl. No. 13/620,910, Aug. 8, 2013, Office Action.  
U.S. Appl. No. 13/620,916, Sep. 6, 2013, Office Action.  
U.S. Appl. No. 13/772,193, Aug. 16, 2013, Office Action.  
AT&T Focus on Communication Feb. 1, 1995, DIANE Publishing Company, InterNIC Directory and Database Services Directory of Directories, pp. 79-85.  
Adorf, H. M.; Resource Discovery on the Internet; May 12, 1995; Elsevier Science Ltd; Vistas in Astronomy, vol. 39, pp. 243-253.  
Sandhu et al., Role-Based Access Control Models, Feb. 1996, IEEE, p. 40.  
Barrows et al., Privacy, Confidentiality, and Electronic Medical Records, Mar./Apr. 1996. Journal of the American Medical Informatics Association, vol. 3, No. 2.  
Johannsen, Thomas; "Representing IP Information in the X.500 Directory"; Mar. 1994; <<http://www.ietf.org/rfc/rfc1608.txt.pdf>>.  
U.S. Appl. No. 12/685,799, Nov. 6, 2013, Office Action.  
U.S. Appl. No. 12/685,799, Jan. 8, 2014, Notice of Allowance.  
U.S. Appl. No. 12/685,799, Feb. 14, 2014, Notice of Allowance.  
U.S. Appl. No. 13/612,390, Jan. 16, 2014, Office Action.  
U.S. Appl. No. 13/612,390, Apr. 16, 2014, Office Action.  
U.S. Appl. No. 13/612,414, Oct. 16, 2013, Office Action.  
U.S. Appl. No. 13/612,414, Dec. 19, 2013, Office Action.  
U.S. Appl. No. 13/612,414, Apr. 11, 2014, Office Action.  
U.S. Appl. No. 13/612,512, Jan. 30, 2014, Office Action.

U.S. Appl. No. 13/620,910, Feb. 13, 2014, Office Action.  
U.S. Appl. No. 13/620,916, Dec. 23, 2013, Office Action.  
U.S. Appl. No. 13/772,193, Feb. 11, 2014, Office Action.  
U.S. Appl. No. 13/620,916, Jun. 9, 2014, Office Action.  
U.S. Appl. No. 13/772,193, Sep. 9, 2014, Office Action.  
U.S. Appl. No. 12/075,581, Apr. 5, 2011, Office Action.  
U.S. Appl. No. 12/075,581, Aug. 31, 2011, Office Action.  
U.S. Appl. No. 12/075,581, Dec. 12, 2011, Office Action.  
U.S. Appl. No. 12/075,581, Apr. 19, 2012, Notice of Allowance.  
U.S. Appl. No. 13/543,846, Jun. 11, 2014, Notice of Allowance.  
U.S. Appl. No. 13/612,390, Oct. 2, 2014, Office Action.  
U.S. Appl. No. 13/612,414, Aug. 26, 2014, Notice of Allowance.  
U.S. Appl. No. 13/612,512, Sep. 12, 2014, Notice of Allowance.  
U.S. Appl. No. 13/620,910, Oct. 9, 2014, Office Action.  
Yuwono, Budi et al.; A World Wide Web Resource Discovery System; Hong Kong University of Science and Technology; Dec. 1995.  
Schwartz, Michael F.; Internet Resource Discovery at the University of Colorado; University of Colorado; Sep. 1993.  
U.S. Appl. No. 13/612,512, Nov. 7, 2014, Notice of Allowance.  
U.S. Appl. No. 13/620,916, Jan. 5, 2015, Office Action.  
U.S. Appl. No. 12/685,799, Jan. 12, 2010, Virdy.  
U.S. Appl. No. 13/620,910, Sep. 15, 2012, Virdy.  
U.S. Appl. No. 13/620,916, Sep. 15, 2012, Virdy.  
U.S. Appl. No. 13/612,414, Sep. 12, 2012, Virdy.  
U.S. Appl. No. 13/612,390, Sep. 12, 2012, Virdy.  
U.S. Appl. No. 13/612,512, Sep. 12, 2012, Virdy.  
U.S. Appl. No. 13/772,191, Feb. 20, 2013, Virdy.

\* cited by examiner

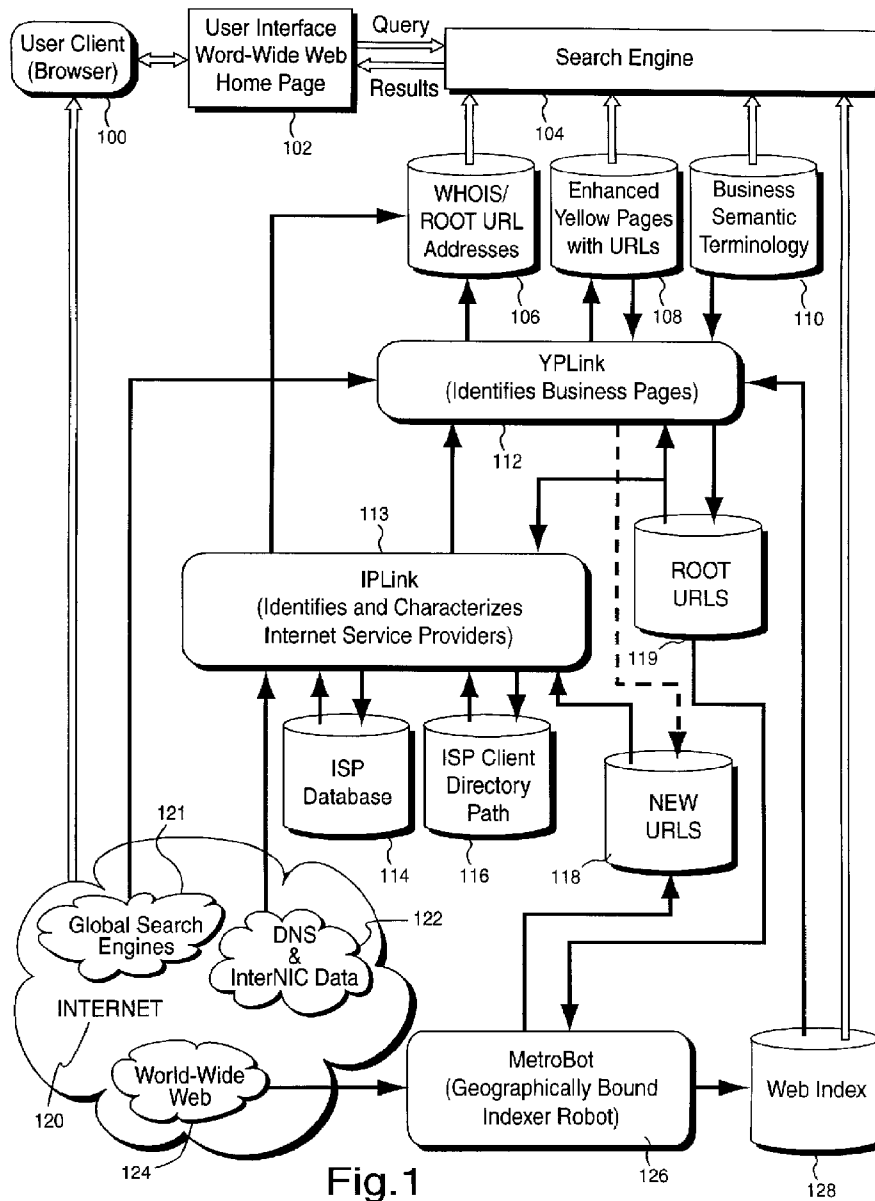
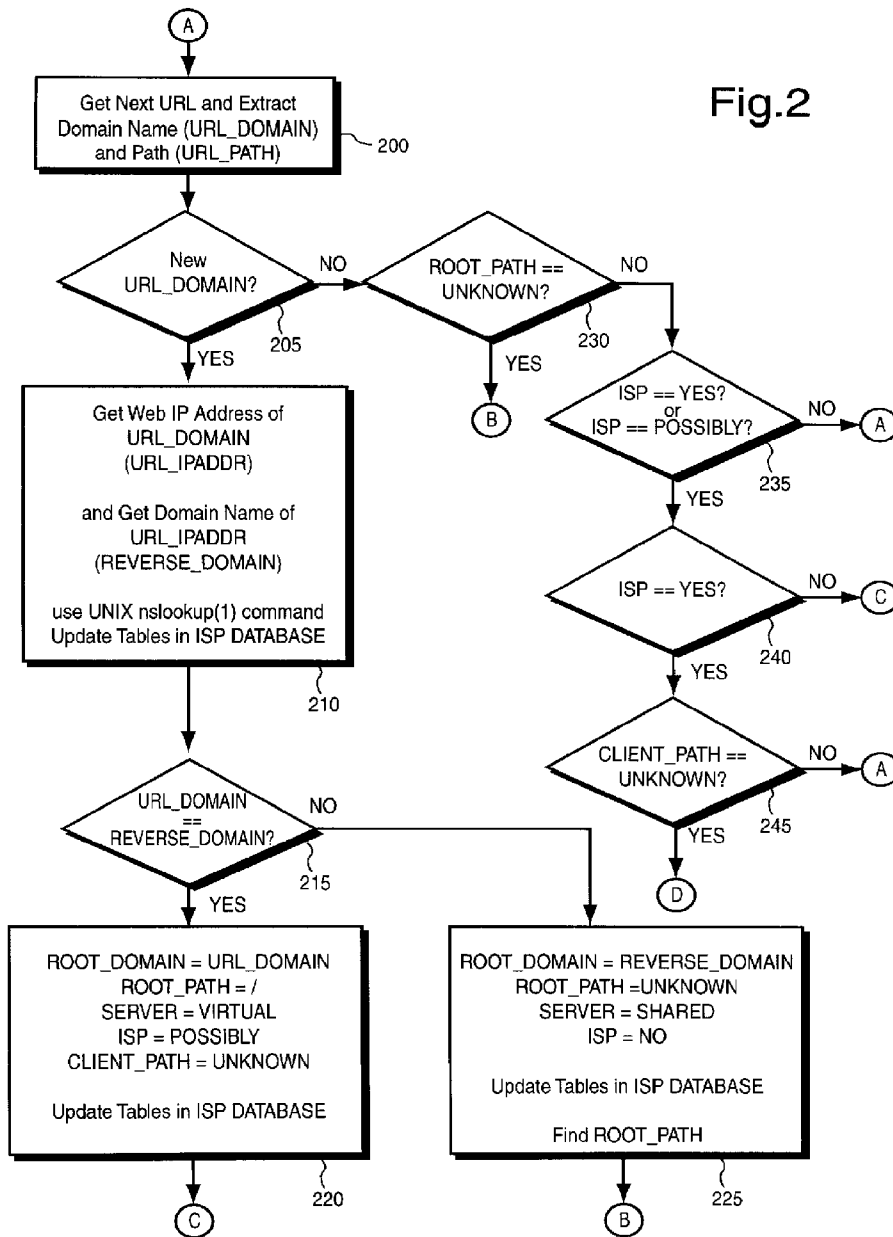


Fig. 1

Fig.2



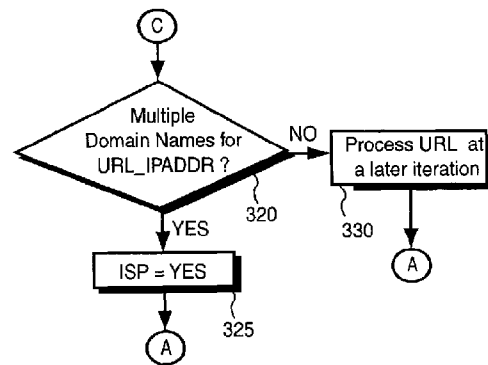
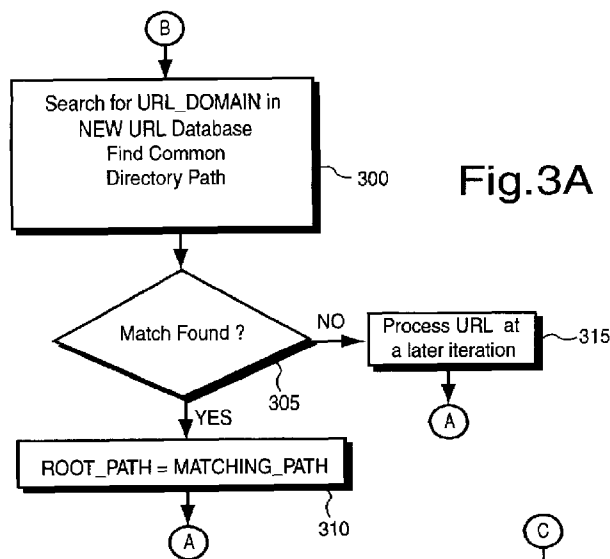


Fig. 3B


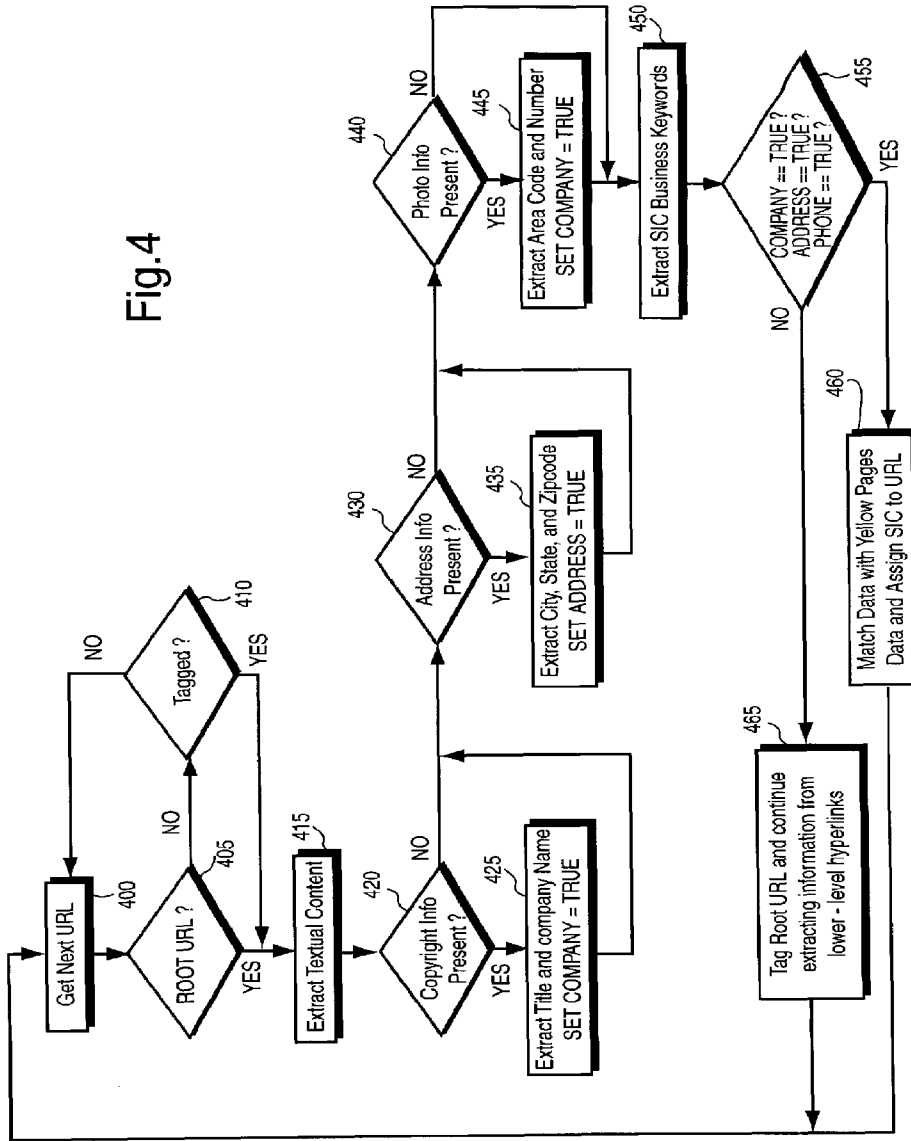


Fig. 4



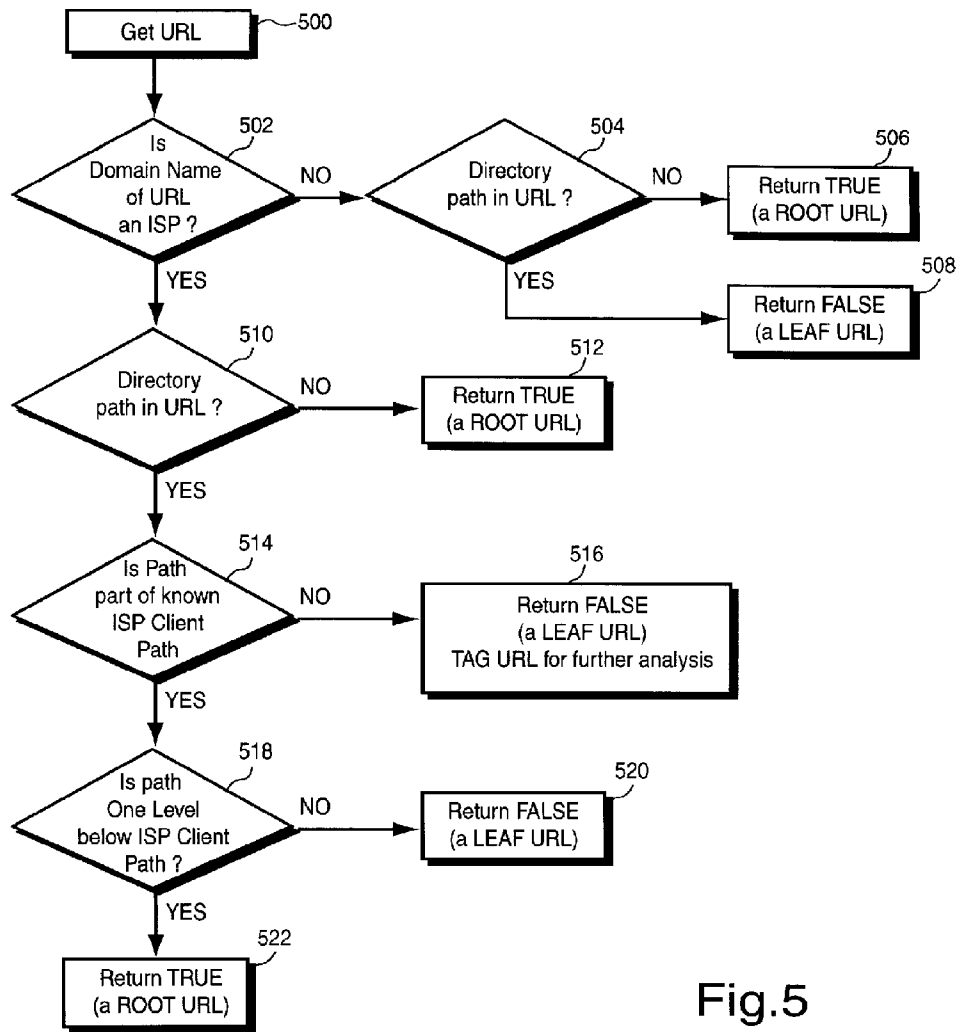


Fig.5



```
<HTML>
<HEAD>
<TITLE>Company Name's Home Page</TITLE>

Company Name
Street Address
City, State Zip
Phone: ###-###-###
Fax: ###-###-###

Copyright (C) 1996 Company Name. All rights reserved.
```

Fig.6

## SYSTEM AND METHOD FOR INDEXING DOCUMENTS ON THE WORLD-WIDE WEB

This application is a continuation of U.S. application Ser. No. 12/685,799 filed on Jan. 12, 2010, which is a continuation of U.S. patent application Ser. No. 12/054,130, filed Mar. 24, 2008, which is a continuation of U.S. patent application Ser. No. 10/705,822, filed Nov. 13, 2003, now U.S. Pat. No. 7,349,892, which is a continuation of U.S. patent application Ser. No. 09/500,471, filed Feb. 9, 2000, now U.S. Pat. No. 6,691,105, which is a continuation of U.S. patent application Ser. No. 08/844,522, filed Apr. 18, 1997, now U.S. Pat. No. 6,148,289, which claims the benefit of U.S. Provisional Application Ser. No. 60/017,548 filed May 10, 1996, the disclosures of which are all incorporated herein by reference in their entireties.

### BACKGROUND OF THE INVENTION

The present invention generally relates to a resource discovery system and method for facilitating local commerce on the World-Wide Web and for reducing search time by accurately isolating information for end-users. For example, distinguishing and classifying business pages on the Web by business categories using the Standard Industrial Classification (SIC) codes is achieved through an automatic iterative process which effectively localizes the Web.

### DESCRIPTION OF THE RELATED ART

Resource discovery systems have been widely studied and deployed to collect and index textual content contained on the World-Wide Web. However, as the volume of accessible information continues to grow, it becomes increasingly difficult to index and locate relevant information. Moreover, global flat file indexes become less useful as the information space grows causing user queries to match too much information.

Leading organizations are attempting to classify and organize all of Web space in some manner. The most notable example is Yahoo, Inc. which manually categorizes Web sites under fourteen broad headings and 20,000 different sub-headings. Still others are using advanced information retrieval and mathematical techniques to automatically bring order out of chaos on the Web.

Solutions to solve this information overload problem have been addressed by C. Mic Bowman et al. using Harvest: A Scalable, Customizable Resource Discovery and Access System. Harvest supports resource discovery through topic-specific content indexing made possible by a very efficient distributed information gathering architecture. However, these topic specific brokers require manual construction and they are geared more for academic and scientific research than commercial applications.

Cornell's SMART engine developed by Gerard Salton uses a thesaurus to automatically expand a user's search and capture more documents. Individual, Inc. uses this system to sift through vast amounts of textual data from news sources by filtering, capturing, and ranking articles and documents based on news industry classification.

The latest attempts for automated topic-specific indexing include the Excite, Inc. search engine which uses statistical techniques to build a self-organizing classification scheme. Excite Inc.'s implementation is based on a modification of the popular inverted word indexing technique which takes into account concepts (i.e., synonymy and homonymy) and analyzes words that frequently occur together. Oracle has devel-

oped a system called ConText to automatically classify documents under a nine-level hierarchy that identifies a quarter-million different concepts by understanding the written English language. ConText analyzes a document and then decides which of the concepts best describe the document's topic.

The systems described above all attempt to organize the vast amounts of data residing on the Web. However, these mathematical information retrieval techniques for classifying documents only work when the message of a document is directly correlated to the words it contains. Attempts to isolate documents by regions or to separate business content from personal content in an automated fashion is not addressed by any conventional system or structure.

### SUMMARY OF THE INVENTION

It is therefore an object of the present invention to provide a method and system for overcoming the above-mentioned problems of the conventional methods and techniques.

The invention is based on a heuristic algorithm which exploits common Web page design principles. The key challenge is to ascertain the owner of a Web page through an iterative process. Knowing the owner of a Web page helps identify the nature of the content business or personal which, in turn, helps identify the geographic location.

In a first aspect of the invention, a method of classifying a source publishing a document on a portion of a network, includes steps of electronically receiving a document, based on the document, determining a source which published the document, and assigning a code to the document based on whether data associated with the document published by the source matches with data contained in a database.

In a second aspect, a search engine is provided for use on a network for distinguishing between business web pages and personal web pages. The search engine includes a mechanism for parsing the content of a hyper-text markup language (HTML) at a web address and searching for criteria contained therein, a mechanism for analyzing a uniform resources locator (URL) of the web address to determine characteristics thereof of a web page at the web address, a mechanism for determining whether the criteria match with data contained in a database, and a mechanism for cross-referencing a match, determined by the determining mechanism, to a second database, to classify a source which published the web page.

### BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, aspects and advantages will be better understood from the following detailed description of a preferred embodiment of the invention with reference to the drawings, in which:

FIG. 1 shows the process flow diagram of a geographically bound resource discovery system including three main components of the invention (sometime referred to below as "MetroSearch") identified as MetroBot, IPLink, and YPLink;

FIG. 2 depicts the IPLink flow chart, the process for identifying ISPs and Client Directory Paths;

FIGS. 3A-3C are sub-processes of the IPLink flow chart shown in FIG. 2;

FIG. 4 depicts the flow chart of YPLink for identifying business pages;

FIG. 5 is a flow diagram for determining if a given uniform resources locator (URL) is a Root URL or a Leaf URL; and FIG. 6 is a template of a typical business home page.

DETAILED DESCRIPTION OF A PREFERRED  
EMBODIMENT OF THE INVENTION

Referring now to the drawings, and more particularly to FIG. 1, there is shown the general arrangement of a preferred embodiment according to the present invention.

The underlying insight behind the invention is that individuals and organizations responsible for the design, creation, and maintenance of their home page generally follow some basic unwritten rules. These rules can be exploited to automatically identify the owner of the home page with a high probability of success. Once the owner of the home page is determined, an SIC code is assigned to it by looking up the owner in a Yellow Pages database. If a matching entry exists, then the owner is a business, otherwise the owner is deemed to be an individual with a personal home page.

FIG. 1 shows a preferred architecture for implementing a geographically bound resource discovery system. The main components of interest are MetroBot 126, IPLink 113, and YPLink 112.

The World-Wide Web ("the Web") 124 is based on a client-server architecture. The Web is the graphical, multimedia portion of the Internet 20. The client side program is a Web browser 100 and the server side is a computer running the HTTPD program 102. The Web server is accessed through the Internet by specifying a Uniform Resource Locator (URL). User-entered queries are sent to a back-end processor or search engine 104 which gathers results from various databases 106, 108, 110, and 128, and formats the request and presents them back to the user.

MetroBot 126 is an indexer robot which traverses hyperlinks in HTML documents and indexes the content into a searchable Web index database 128. These hyperlinks or URLs point to other Web pages making it possible to recursively traverse large portions of the Web from a single, well-chosen URL (seed URL). MetroBot begins its traversal from known Root URL 119 such as the home page of a local service provider (SP), such as an internet service provider (ISP). New links that are discovered are stored in New URLs database 118. These links are processed by IPLink 113 and YPLink 112 to extract new Root URLs at which point the whole process repeats itself. Furthermore, YPLink periodically supplements its New URL list by querying global search engines 121 using strategic keywords (e.g., regional city, county, state names, zip codes, and industry specific terms).

The first level of localization is achieved by limiting URLs to registered domain names 106. IPLink extracts domain names from the New URL database and then queries the InterNIC database 122 where records of registered domain names containing company name, contact, street address, and Internet Protocol (IP) addresses are kept. This InterNIC database can be accessed through the Unix whois(1) command. YPLink merges the InterNIC address database 108 with the Yellow Pages data 110. This process is described in detail below.

The next level of localization is more complex since most businesses do not have their own registered domain name. Instead, they have their home page hosted on local SPs (or ISPs) or Online Service Providers (OSPs) Web Servers.

The first step in solving this problem is for IPLink 113 to characterize URLs by their IP addresses. FIGS. 2 and 3A-3C show the IPLink flow logic. IPLink identifies the following attributes based on the IP addresses of New URLs:

- True/Virtual Web Servers vs. Shared Web Servers.
- ISP vs. Non-ISP hosts.
- Root Domain of URLs.
- Root Path of URLs.
- Client Directory Paths if host is an ISP.

A new URL is retrieved from the New URL database 200 and is parsed into the domain name and directory path portions. If it is a new domain 205, then its Web IP address (i.e., www.domain.name) is retrieved using the Internet Domain Name Service 122. The Unix nslookup(1) utility 210 returns an IP address given a domain name. The corresponding IP address is stored in the ISP database 114. A reverse lookup 210 of the Web IP address is also performed to determine 215 if the given URL is hosted on a true (or virtual) Web server 220 or a shared Web server 225. A domain name with its own unique Web IP address indicates a true or virtual Web server (non-ISP host). Multiple domain names for a single Web IP address indicates a shared Web server (ISP host).

The official domain name (Root Domain) 220 and 225 for the IP address is the domain name of the ISP (master/slave name server information returned by whois(1) can also be used to accurately identify the ISP if the Root Domain does not correspond to the ISP). Root Domain is only used for displaying URL information on search results not for further processing.

Turning to FIG. 3A, for shared servers 225, the Root Path is determined by searching 300 for the given domain name in the New URL database 118 and finding common directory paths 305. If no match is found 315, the URL will automatically be processed at a later iteration 230, otherwise the Root Path is set to the matching path 310. Turning to FIG. 3B, for virtual servers 220, the Root Path is simply the root directory ('/'). These servers may or may not be ISPs. If multiple domain names exist for the given IP address 320, then it is classified as an ISP 325, otherwise it is processed at a later iteration 330, 235 and 240. It is possible for organizations to become ISPs in the future by simply adding/hosting new domain names on their existing Web servers.

The directory path where the ISP stores its customers Web pages is called the ISP Client Directory Path 116. This data is initially created manually for a few local ISPs (seed ISPs). This path is identified automatically 335 by searching for the given domain name in the Root URL database 119 and finding common directory paths 340, as shown in FIG. 3C. If no match is found 350, then it is processed at a later iteration 245. Matching paths 345 point to the ISPs Client Directory Path. This process improves over subsequent iterations when enough data is gathered and patterns can be recognized from a large set of ISP Web Servers.

IPLink encompasses the first phase of identifying and characterizing IP addresses. The next phase is to automatically identify businesses hosted on ISP Web servers.

FIG. 4 shows the YPLink flow chart. YPLink determines if a Web page belongs to a business or an individual. YPLink takes its input, a URL, from IPLink. FIG. 4 shows the flow diagram for the YPLink process. The first step after retrieving a URL 400 is determining if it is a "Root URL" or a "Leaf URL" 405.

A Root URL is the entry point for an organization's or individual's home page on the World-Wide Web. A Root URL may or may not be the same as the Home page. Leaf URLs, on the other hand, are links below an organization's Root URL.

Four factors are considered in determining a Root URL:

1. Is the URL hosted on a Service Provider's Web Server?
2. Is the URL on a virtual Web Server?
3. Does the URL contain a directory path?
4. Is the directory path a known Service Provider's Client Directory?

IPLink determines the SP Client Directory Path as described above. The ISP database 114 contains information about Client Directories for various ISPs.

FIG. 5 shows the Root URL flow logic. A given URL is retrieved 500 and parsed into two components: domain name and directory path. The domain name is analyzed to see if it is an ISP 502. If multiple IP addresses are associated with the

domain name, then the domain name is an ISP. If the domain name is not an ISP, then the directory path component is checked **504**. A missing directory path signifies a Root URL **506**, otherwise it is a Leaf URL **508**.

If the domain name is an ISP **510**, then it is also a Root URL if no directory path exists **512**. If a directory path exists **514**, then the path is compared to a list of known ISP Client Directory paths. No match **516** indicates a Leaf URL, otherwise the directory path level is analyzed **518** for final Root URL determination. If the path is one directory level below the Client Directory path then it is a Root URL **522**, otherwise it is a Leaf URL **520**.

After a URL is determined to be Root URL, then the home page it points to is analyzed **415** to see if it follows some basic guidelines. A typical home page layout is illustrated in FIG. **6**. Other than following HTML requirements, there is no rule or standards for the layout of textual content. The key pieces of information required to ascertain the owner of a Web page are 1) company name, 2) zip code, and 3) telephone number. These three pieces of information do not have to exist in the Root URL. They can reside anywhere among various Leaf URLs beneath a Root URL. In many cases, this information is stored in a file called about .html. However, the same information could be stored in other, similarly named files, as would be known to those skilled in the art taking the present specification as a whole. The process described below extracts this information automatically and assigns it to the Root URL being analyzed.

The company's name is usually included in the HTML TITLE tag **600**. However, the company's name could be included in other locations, as would be known to those ordinarily skilled in the art within the purview of the present specification. The layout of the address, if present, usually is in a standard recognizable format **602**. Most businesses also tend to include copyright notices near the bottom of their documents. A string search for "copyright", "©", and "&#169;" is performed near the bottom **604** of the home page. The company name usually appears near the copyright notice. A match of the organization or individual's name in the copyright field **420** and the TITLE field **425** provides the first indication of the owner of the home page. If no match is found, then the URL is tagged for further analysis during the next iteration.

The next step is to analyze the URL for address **430** information. Addresses have an easily identifiable format. In the U.S., the format is the city name followed by a comma and then followed by the full state name or abbreviation and finally a five or nine digit zip code. However, other common formats/methods also are possible and would be known to those ordinarily skilled in this art field to locate the zip code. This string is parsed in the HTML file after stripping all tags **435**. The only information required is the 5-digit zip code since the city and state can be determined by this field alone. YPLink stores addresses associated with Root URLs and domain names in an address database **106**.

If a phone format field is present then it is also extracted and stored **440**. U.S. phone field is a 10-digit field where the first three digits representing the area code are optionally enclosed in parentheses or separated by a dash, space, or a period, and then followed by a 7-digit number which is separated by a dash, space, or a period after the third digit **445**. Other similar methods of identifying a phone number are known to those ordinarily skilled in the art.

The pair consisting of the company name and zip code are usually enough to identify a business **455**. A query is constructed using this pair and sent to a Yellow Pages database server. This database is indexed by business names and zip

codes. If a single match is found, then the resulting SIC code is assigned to the corresponding Root URL **460**. If multiple entries are matched, then the phone field is also included in the query to assure that only a single entry is retrieved. If no match is found, then the URL is tagged **465** for further analysis of lower-level hyperlinks during the next iteration. The matching data is stored in an enhanced Yellow Pages database **108**.

If no match is found at any level, then the page is tagged **450** as a personal page with an SIC code assigned according to the closest match based on the Business Semantic Terminology database **110**. This database is a proprietary thesauri of keywords relating business categories in the Yellow Pages and other emerging industries such as Internet technology to extended SIC codes.

While the invention has been described in terms of a single preferred embodiment, those skilled in the art will recognize that the invention can be practiced with modification within the spirit and scope of the appended claims.

For example, while the invention above has been described primarily in terms of (e.g., implemented in) a software process and a system employing software and hardware, the invention could also be implemented with hardware as would be known by one of ordinary skill in the art taking the present specification as a whole.

What is claimed is:

**1.** A method comprising:

maintaining a URL database comprising a plurality of URLs;

retrieving, without user intervention, a first document associated with a first URL from the plurality of URLs;

parsing, using at least one processor, the first document to identify one or more additional URLs within the first document;

comparing the one or more additional URLs to the URL database to determine if the URL database includes the one or more additional URLs;

if it is determined that the URL database does not include the one or more additional URLs, updating the URL database to include the one or more additional URLs;

determining if the one or more additional URLs are root URLs based on a set of predefined rules;

when it is determined that the one or more additional URLs are root URLs, classifying the one or more additional URLs as root URLs in the URL database and updating the URL database to add the classification that the one or more additional URLs are root URLs;

querying a third-party database for additional information associated with the one or more additional URLs;

updating the URL database to include the additional information associated with the one or more additional URLs; and

repeating the steps of retrieving, parsing, comparing, and updating for each of the one or more additional URLs.

**2.** The method of claim **1**, further comprising determining whether the first URL is a root URL.

**3.** The method of claim **2**, further comprising obtaining geographical information associated with the first URL from the third-party database.

**4.** The method of claim **1**, wherein determining if the one or more additional URLs are root URLs based on the set of predefined rules comprises:

parsing the one or more additional URLs into a domain name and a directory path; and

determining that the domain name is an ISP.

5. The method of claim 1, wherein determining if the one or more additional URLs are root URLs based on a set of predefined rules comprises:

- parsing the one or more additional URLs into a domain name and a directory path;
- determining that multiple IP addresses are associated with the domain name; and
- identifying that the directory path is one directory level below the domain name.

6. The method of claim 1, wherein repeating the steps of retrieving, parsing, comparing, and updating for each of the one or more additional URLs is prioritized to favor the one or more additional URLs classified as root URLs.

7. The method of claim 1, further comprising, if it is determined that the one or more additional URLs are not root URLs, classifying the one or more additional URLs as leaf URLs in the URL database.

8. The method of claim 1, further comprising extracting a domain name for each of the one or more additional URLs.

9. The method of claim 8, further comprising querying a domain name database with the extracted domain names from the one or more additional URLs, the domain name database maintaining information associated with registered domain names.

10. The method of claim 9, further comprising receiving, from the domain name database, information associated with the extracted domain names.

11. The method of claim 10, further comprising updating the updating the URL database to include the information associated with the extracted domain names.

12. The method of claim 11, further comprising searching the updated URL database using a search term.

13. The method of claim 12, further comprising identifying resultant URLs based on matching the search term with the information associated with the one or more additional URLs within the URL database.

14. The method of claim 10, wherein the information associated with the extracted domain names comprises geographic location information.

15. The method of claim 14, further comprising searching the information associated with the extracted domain names using a search term that includes a geographic location component.

16. The method of claim 1, wherein the URL database comprises a URL index.

17. A method comprising:

maintaining a URL database comprising one or more URLs;

retrieving one or more documents associated with the one or more URLs;

parsing, using at least one processor, the one or more documents to identify a first set of new URLs within the one or more documents;

determining if one or more URLs in the first set of URLs are root URLs based on a set of predefined rules;

if it is determined that one or more URLs in the first set of URLs are root URLs, classifying the one or more URLs in the first set of URLs as root URLs in the URL database and updating the URL database to include the identified first set of new URLs including the one or more URLs classified as root URLs;

querying, without user intervention, a third-party search engine using one or more keywords to identify a second set of new URLs; and

updating the URL database to include the identified second set of new URLs.

18. The method of claim 17, further comprising obtaining geographical information associated with the first URL from the third-party database.

19. The method of claim 17, wherein updating the URL database to include the identified first set of new URLs comprises comparing the first set of new URLs to the one or more URLs to determine if the URL database includes each of the first set of new URLs.

20. The method of claim 19, wherein updating the URL database to include the identified first set of new URLs further comprises, if it is determined that the URL database does not include one or more of the first set of new URLs, adding the one or more of the first set of new URLs to the URL database.

21. The method of claim 17, further comprising associating, within the URL database, the second set of new URLs with the one or more keywords used to identify the second set of new URLs.

22. The method of claim 21, wherein the one or more keywords comprise a geographic identifier.

23. The method of claim 21, wherein the one or more keywords comprise an industry identifier.

24. The method of claim 17, further comprising extracting a domain name for each of the first set of new URLs.

25. The method of claim 24, further comprising querying a domain name database with the extracted domain names from the first set of new URLs, the domain name database maintaining information associated with registered domain names.

26. The method of claim 25, further comprising receiving, from the domain name database, information associated with the extracted domain names from the first set of new URLs.

27. The method of claim 26, further comprising associating, within the URL database, the first set of new URLs with the received information associated with the extracted domain names from the first set of new URLs.

28. The method of claim 27, wherein the information associated with the extracted domain names comprises geographic information.

29. The method of claim 17, wherein the URL database comprises a URL index.

30. A system comprising:

at least one processor; and

a non-transitory computer-readable storage medium storing instructions thereon that, when executed by the at least one processor, cause the at least one processor to: maintain, on a non-transitory computer readable medium, a URL database comprising one or more URLs;

query, without user intervention, a first third-party database for information associated with the one or more URLs;

query, without user intervention, a second third-party database for information associated with the one or more URLs;

merge the information associated with the one or more URLs from the first third-party database and the information associated with the one or more URLs from the second third-party database;

determining if the one or more URLs are root URLs based on a set of predefined rules;

if it is determined that the one or more URLs are root URLs, classifying the one or more URLs as root URLs in the URL database and updating the URL database to add the classification that the one or more URLs are root URLs; and

updating the URL database to include the merged information from the first third-party database and the second third-party database.

31. The system of claim 30, wherein the additional information associated with the one or more URLs comprises geographical information.

32. The system of claim 30, wherein the first third-party database is a domain name database. 5

33. The system of claim 32, wherein the second third-party database is a business listing database.

34. The system of claim 33, further comprising instructions that, when executed by the at least one processor, cause the at least one processor to receive a search term from a user. 10

35. The system of claim 34, further comprising instructions that, when executed by the at least one processor, cause the at least one processor to query the URL database based on the search term.

36. The system of claim 35, further comprising instructions 15 that, when executed by the at least one processor causes the at least one processor to identify one or more resultant URLs having merged information that correlates to the search term.

37. The system of claim 36, further comprising instructions 20 that, when executed by the at least one processor, cause the at least one processor to provide the one or more resultant URLs to the user.

38. The system of claim 30, wherein the URL database comprises a URL index.

\* \* \* \* \*